CRD

Computational Research Division Report

Berkeley Lab

Office of Science
U.S. Department of Energy

AUGUST 2008

Technology for Speeding Up Searches of Large Databases Wins R&D 100 Award

FastBit, an indexing technology that allows users to search massive datasets up to 40 times faster than the best commercially available tools, has been recognized with a 2008 R&D 100 Award.

R&D Magazine will present the award to FastBit developers Kesheng "John" Wu, Arie Shoshani, Ekow Otoo, and Kurt Stockinger of the Scientific Data Management Group in Berkeley Lab's Computational Research Division, at a special October ceremony in Chicago, Ill. (Stockinger has since left the group to accept a position in Switzerland.

According to Wu, who led the project's

continued on page 3

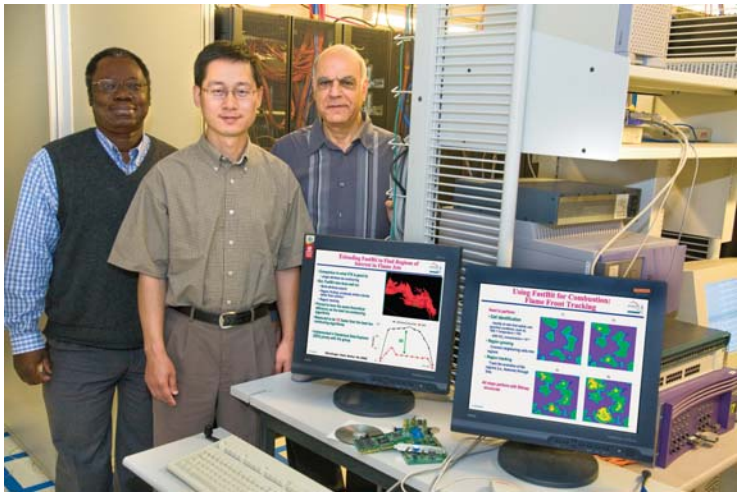

Kesheng "John" Wu (center) leads Fastbit to the top of the 2008 R&D100. Award winning team members include, Ekow Otoo (left), Arie Shoshani (right), and Kurt Stockinger (not pictured).

Juan Meza Wins 2008 Blackwell-Tapia Prize and SACNAS Distinguished Scientist Award

Juan C. Meza, Department Head and Senior Scientist for the High Performance Computing Research Department in the Computational Research Division at Lawrence Berkeley National Laboratory, has been named recipient of two awards: the Blackwell-Tapia Prize and the SACNAS Distinguished Scientist Award.

Blackwell-Tapia Prize

The Blackwell-Tapia Prize is awarded every second year in honor of the legacy of David H. Blackwell and Richard A. Tapia, two distinguished mathematical scientists who have been inspirations to more

continued on page 4

Steve Cotter Named New Head of ESnet

Steve Cotter will be supporting science around the world as the new head of ESnet, the Department of Energy's high-speed network, beginning Friday, August 29.

For more than a decade, Cotter has designed and deployed networks, for both research and commercial use, at home and abroad. Most recently, he served as Google's network deployment manager for Europe, the Middle East, and Africa. At ESnet, he succeeds Bill Johnston, who is retiring from Berkeley Lab after more than 35 years.

"Steve Cotter brings to this job exactly the right combination of experience in all aspects of network design, development, deployment and operation," said LBNL Associate Laboratory Director Horst Simon. "On top of that, Steve is already familiar with ESnet's staff and operations and has strong connections to the research network community. We are thrilled by his decision to lead ESnet."

Before joining Google in 2007, Cotter worked for Internet2, a high performance network serving more than 300 institutions in the research and education community in the U.S. Since 2006, ESnet and Internet2 have worked as partners in building ESnet's next-generation infrastructure.

"While at Internet2, I worked closely with ESnet on the design of their new infrastructure and became familiar with their research community and staff, a very impressive group," Cotter said. "Bill Johnston is leaving some big shoes to fill, but it's an exciting challenge and I welcome it. The services ESnet provides

continued on page 2

Cotter *continued from page 1*

are not just good for the research community, but also the public by supporting important research in areas like global climate change and new and renewable energy sources."

ESnet, or the Energy Sciences Network, is managed by Lawrence Berkeley National Laboratory for the Department of Energy. ESnet provides direct connections to more than 40 DOE sites, as well as fast interconnections to more than 100 other networks. Funded principally by DOE's Office of Science, ESnet services allow scientists to make effective use of unique DOE research facilities and computing resources, independent of time and geographic location.

Based on his experience in the commercial sector, Cotter said he sees that ESnet is not alone in pushing the envelope of advanced networking and believes the entire community can benefit from continuing ESnet's strong record of building alliances with other organizations to advance the state of networking. Key areas include transferring massive sets of scientific data and supporting "cloud" computing.

Built for speed: When Steve Cotter is not heading ESnet, he's zipping through the French Alps on bike.

"Google is an incredible place to work and I learned a lot, but ESnet is doing some amazing things as well," Cotter said. "This is an opportunity to be part of a great team and to make an impact."

Cotter brings a broad range of experience with other leading firms to his new role. While serving as Internet2's Deputy Operations Officer, their close collaboration with ESnet's OSCARs project and others in the research community resulted in the rollout of Internet2's Dynamic Circuit Network – delivering on NSF's Cyberinfrastructure vision that the network should be available as a schedulable, on-demand resource like supercomputers are today. Cotter said "I was also the CEO of Internet2's subsidiary FiberCo — brokering dark fiber acquisitions and providing consulting/project management services to the higher education and research community."

Before that, Cotter was assigned from Internet2 to National LambdaRail, where he led the build-out and design of their DWDM (dense wave division multiplexing) and IP network. He got his start in optical network technologies while an engineer at Cisco Systems and also served as a Director in SBC Global Communications' (now AT&T) engineering and construction departments. He is well known in the networking community, having been a speaker at networking meetings and conferences in the U.S. and Europe.

Although he has extensive experience in industry, Cotter said his approach to management was shaped during his eight years as a combat helicopter pilot in the U.S. Marine Corps. Cotter led missions in both the former Yugoslavia and Somalia and Liberia in Africa, facing everything from hostile fire to the most difficult weather conditions.

"You learn to focus on the mission at hand, to push forward no matter what obstacles you come across, and to take care of those you are responsible for," Cotter said. "You rely on your people to do their job, to do it well and to work as a team. I've carried over that philosophy of having faith in those who work for you by letting them do their job and getting out of their way."

Raised in New Jersey, Cotter attended the U.S. Naval Academy, where he earned his bachelor's degree in aeronautical engineering. He earned an MBA from Boston University and has completed about 75 percent of his coursework toward a master's degree in Information Economics, Management and Policy at the University of Michigan.

Cotter and his wife, Pallas, have two young sons and live in Davis, Calif.

> "You rely on your people to do their job, to do it well and to work as a team. I've carried over that philosophy of having faith in those who work for you by letting them do their job and getting out of their way." – Steve Cotter

R&D 100 Award *continued from page 1*

development, FastBit was originally designed for particle physicists who need to sort through billions of data records just to find tens or hundreds of key pieces of information. For example, the technology has played a crucial role in a high energy physics experiment called STAR, where colliding particles generate billions of subatomic events, but only a few hundred of these events have the most distinctive signatures of the new state of matter called quark-gluon plasma (QGP). FastBit helped the researchers quickly identify evidence of QGPs, which was a key objective of the STAR experiment.

Wu notes that the key to FastBit's success is an innovative bitmap compression method called the Word Aligned Hybrid (WAH) compression. Computer scientists have long known that bitmap indexing methods are the most efficient for accelerating searching operations on large datasets that do not change over time, which is the case for most science experiments.

Commonly used commercial indexing methods, such as B-tree and similar tree-based indexes, are inefficient for scientific research because they cannot locate thousands of records in analyses of large datasets. Instead, these technologies are optimized for finding a small number of specific records, such as one's bank account information.

Additionally, B-tree indexes inflate the volume of the data by a factor of three or four, which can be extremely cumbersome when dealing with terabyte-size datasets. Meanwhile, many tree-based indexing techniques are specialized to the task at hand – meaning they are tailored for searching regions on a single variable, such as finding all hot regions in a climate database – but do not work well when multiple variables are searched together, such as finding hot regions where wind velocity is high.

Although bitmap indexes proved to be better, they also had several shortcomings. These indexes could only search limited types of data, primarily those variables with a relatively small number of possible values, such as the sex of a customer or the state in which the customer lives. Another limitation was size – bitmap indexes were too large to be stored in computer memory, so to answer a query, the relevant part of the index had to be read into the computer's memory before the necessary computation could be performed. This significantly slowed down a search because the time needed to read parts of the index into the computer's memory tends to be much greater than the time needed to perform the actual query.

> "In short, FastBit contains significant innovations that are well-recognized and have wide impacts in science, technology and education." – John Wu

The FastBit team overcame these issues by inventing a compression method that was demonstrated to be 10 times faster than its widely used commercial counterpart. In addition to being compute-efficient, it is also effective in reducing the index size. The size of a FastBit index is on average one-third the size of the original data, which is about one-tenth the size of the most widely used B-tree index.

FastBit also significantly improved the speed of searching operations with a number of techniques, including a vertical data organization, an innovative bitmap compression technique, and several new bitmap encoding methods.

"Since we began developing FastBit to help physicists search massive datasets resulting from large-scale experiments, we've seen an explosion of data in many other fields, too. We released FastBit under an open-source license in 2007, [and] it has attracted a lot of interest in new areas, such as network traffic data analysis, and business applications, drug discovery and web content delivery," said Wu.

"In short, FastBit contains significant innovations that are well recognized and have wide impacts in science, technology and education."

The R&D100 Award for the public release of FastBit is the latest in a series of honors from the research community. FastBit received an award from the 2005 International Supercomputing Conference for enabling Grid-based analysis of high energy physics data. In the High Performance Analytics Challenge at the Supercomputing 2005 conference, Fastbit developers' work also received an honorable mention for its contributions to network traffic analysis. In addition, FastBit publications have appeared in journals and conferences spanning both the database research field and visualization field, including ACM Transactions on Database Systems, Very Large Database conference and IEEE Visualization conference. FastBit has played crucial roles in two Ph.D. theses from the University of California Berkeley and the University of Illinois Urbana-Champaign.

The project was supported by the Department of Energy's Scientific Discovery through Advanced Computing program, which develops software for accelerating computational science.

For more information about FastBit, go to: http://sdm.lbl.gov/fastbit/.

Blackwell-Tapia Prize *continued from page 1*

than a generation of African American, Latino/Latina, and Native American students and professionals in the mathematical sciences. It recognizes a mathematical scientist who has contributed and continues to contribute significantly to research in his or her field of expertise, and who has served as a role model for mathematical scientists and students from under-represented minority groups or contributed in other significant ways to addressing the problem of the under-representation of minorities in mathematics. The current award citation said:

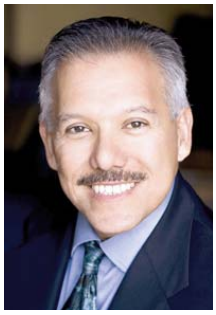

Juan Meza

> Dr. Meza has an exceptionally distinguished record as a mathematical scientist, an accomplished and effective head of a large department doing cutting-edge explorations in the computational sciences, computational mathematics, and future technologies, and a role model and active advocate for others from groups under-represented in the mathematical sciences. As a mathematician, his current research focuses on nonlinear optimization with an emphasis on methods for parallel computing, and he has also worked on various scientific and engineering applications including scalable methods for nanoscience, power grid reliability, molecular conformation problems, optimal design of chemical vapor deposition furnaces, and semiconductor device modeling. He is a much sought after speaker, both nationally and internationally, on topics ranging from his own research, through major invited talks on the importance of diversity such as his presentation as the 2008 Marjorie Lee Browne Colloquium Speaker for the University of Michigan's Martin Luther King, Jr. celebration, and advice important to young mathematicians-in-the-making such as his presentations to student groups on how they can be effective speakers and presenters themselves. His record of service to communities under-represented in mathematics includes co-organizing the Tenth CAARMS Conference in Berkeley, chairing the Mathematical Sciences Research Institute (MSRI) Human Resources Advisory Committee, co-chairing the annual Diversity Day workshops of the Society for Industrial and Applied Mathematics, and many other activities too numerous to mention here; however, they regularly extend from serving on high-level advisory committees on diversity for major scientific organizations, through rolling up his own sleeves and working directly with early-career mathematics students from under-represented groups, as he did in the 2007 MSRI Undergraduate Program (MSRI-UP).

The prize will be presented at the Fifth Blackwell-Tapia Conference, to be hosted by the Statistical and Applied Mathematical Sciences Institute (SAMSI) in Research Triangle Park, NC on November 14–15, 2008. www.samsi.info/workshops/2008Blackwell-Tapia.shtml

SACNAS Distinguished Scientist Award

The mission of SACNAS (Society for Advancement of Chicanos and Native Americans in Science) is to encourage Chicano/Latino and Native American students to pursue graduate education and obtain the advanced degrees necessary for science research, leadership, and teaching careers at all levels. The SACNAS Distinguished Awards recognize scientific achievement, teaching, and mentorship of underrepresented minority students.

"The individuals to be recognized represent the highest caliber of mentoring and an established record of encouraging minority students to pursue advanced degrees in science, mathematics, and engineering," according to the award announcement, which continued, "[Dr. Meza] leads a group conducting cutting-edge research involving the marriage of new algorithmic techniques and advanced computing platforms. Dr. Meza has also worked on various scientific and engineering applications including scalable methods for nanoscience and power grid reliability. In addition, Dr. Meza is extremely active nationally in addressing the underrepresentation of minorities in mathematics."

The award will be presented by SACNAS president Aaron Velasco, Ph.D., during the opening ceremony of the organization's annual national conference on Thursday, October 9, 2008, at the Salt Palace Convention Center in Salt Lake City, Utah. http://www.sacnas.org/

About CRD Report

CRD Report, which publishes every other month, highlights the cutting-edge research conducted by staff scientists in areas including turbulent combustion, nanomaterials, climate change, distributed computing, high-speed networks, astrophysics, biological data management and visualization. CRD Report Editor Linda Vu can be reached at 510 495-2402 or **LVu@lbl.gov**. Find previous CRD Report articles at http://crd.lbl.gov/html/news/CRDreport.html.